

Concrete Model Checking with Abstract Matching and Refinement

Corina S. Păsăreanu¹, Radek Pelánek^{*2}, and Willem Visser¹

¹ NASA Ames Research Center

² Masaryk University Brno, Czech Republic

Abstract. We propose an abstraction-based model checking method which relies on refinement of an under-approximation of the feasible behaviors of the system under analysis. The method preserves errors to safety properties, since all analyzed behaviors are feasible by definition. The method does not require an abstract transition relation to be generated, but instead executes the concrete transitions while storing abstract versions of the concrete states, as specified by a set of abstraction predicates. For each explored transition, the method checks, with the help of a theorem prover, whether there is any loss of precision introduced by abstraction. The results of these checks are used to decide termination or to refine the abstraction, by generating new abstraction predicates. If the (possibly infinite) concrete system under analysis has a finite bisimulation quotient, then the method is guaranteed to eventually explore an equivalent finite bisimilar structure. We illustrate the application of the approach for checking concurrent programs. We also show how a lightweight variant can be used for efficient software testing.

1 Introduction

Over the last few years, model checking based on abstraction-refinement has become a popular technique for the analysis of software. In particular the abstraction technique of choice is a property preserving over-approximation called predicate abstraction [9] and the refinement removes spurious behavior based on automatically analyzing abstract counter-examples. This approach is often referred to as CEGAR (counter-example guided automated refinement) and forms the basis of some of the most popular software model checkers [1, 2, 12]. Furthermore, a strength of model checking is its ability to automate the detection of subtle errors and to produce traces that exhibit those errors. However, over-approximation based abstraction techniques are not particularly well suited for this, since the detected defects may be spurious due to the over-approximation - hence the need for refinement. We propose an alternative approach based on refinement of under-approximations, which effectively preserves the defect detection ability of model checking in the presence of aggressive abstractions.

* Partially supported by GA ČR grant no. 201/03/0509

The technique uses a novel combination of (explicit state) model checking, predicate abstraction and automated refinement to efficiently analyze increasing portions of the feasible behavior of a system. At each step, either an error is found, we are guaranteed no error exists, or the abstraction is refined. More precisely, the proposed model checking technique traverses the concrete transitions of the system and for each explored concrete state, it stores an abstract version of the state. The abstract state, computed by predicate abstraction, is used to determine whether the model checker’s search should continue or backtrack (if the abstract state has been visited before). This effectively explores an under-approximation of the feasible behavior of the analyzed system. Hence all counter-examples to safety properties are preserved.

Refinement uses weakest precondition calculations to check, with the help of a theorem prover, whether the abstraction introduces any loss of precision with respect to each explored transition. If there is no loss of precision due to abstraction (we say that the abstraction is *exact*) the search stops and we conclude that the property holds. Otherwise, the results from the failed checks are used to refine the abstraction and the whole verification process is repeated anew. In general, the iterative refinement may not terminate. However, if a finite bisimulation quotient [14] exists for the system under analysis, then the proposed approach is guaranteed to eventually explore a finite structure that is bisimilar to the original system.

The technique can also be used in a lightweight manner, without a theorem prover, i.e. the refinement guided by the exactness checks is replaced with refinement based on syntactic substitutions [16] or heuristic refinement. The proposed technique can be used for systematic testing, as it examines increasing portions of the system under analysis. In fact, our method extends existing approaches to testing that use abstraction mappings [10, 23], by adding support for automated abstraction refinement.

To the best of our knowledge, the presented approach is the first predicate abstraction based analysis which focuses on automated refinement of under-approximations with the goal of efficient error detection. We illustrate the application of the approach for checking safety properties in concurrent programs and for testing container implementations.

Comparison with Related Work The most closely related work to ours is that of Grumberg et al. [11] where a refinement of an under-approximation is used to improve analysis of multi-process systems. The procedure in [11] checks models with an increasing set of allowed interleavings of the given processes, starting from a single interleaving. It uses SAT-based bounded model checking for analysis and refinement, whereas here we focus on explicit model checking and predicate abstraction, and we use weakest precondition calculations for abstraction refinement.

Our approach can be contrasted with the work on abstracting modal transition systems (see e.g. [19]) for the verification and refutation of branching time temporal logic properties. An abstract model for such logics distinguishes between *may* transitions, which over-approximate transitions of the concrete

model, and *must* transitions, which under-approximate the concrete transitions. The method presented here explores and generates a structure which is *more precise*, i.e. contains more feasible behaviors, than the model defined by the *must* transitions (for the same abstraction predicates). This is due to the fact that the model checker explores transitions that correspond not only to *must* transitions, but also to *may* transitions that are feasible (see discussion in Section 2).

Moreover, unlike [19] and over-approximation abstraction techniques [1, 2], the under-approximation and refinement approach does not require the a priori construction of the abstract transition relation, which involves exponentially many theorem prover calls (in the number of predicates), regardless of the size of (the reachable portion of) the analyzed system. In our case, the model checker executes concrete transitions and a theorem prover is only used during refinement, to determine whether the abstraction is exact with respect to each executed transition. Every such calculation makes at most two theorem prover calls, and it involves only the *reachable* state space of the system under analysis. Another difference with previous abstraction techniques, is that the refinement process is not guided by the spurious counter-examples (since no spurious behavior is explored). Instead, the refinement is guided by the failed exactness checks for the explored transitions.

In previous work [17], we developed a technique for finding guaranteed feasible counter-examples in abstracted programs. The technique essentially explores an under-approximation defined by the *must* abstract transitions (although the presentation is not formalized in these terms). The work presented here explores an under-approximation which is more precise than the abstract system defined by the *must* transitions. Hence it has a better chance of finding bugs while enabling more aggressive abstraction and therefore more state space reduction.

Model-driven software verification [13] advocates the use of abstraction mappings during concrete model checking in a way similar to what we present here. The CMC model checking tool [15] also attempts to store state information in memory using aggressive compressing techniques (which can be seen as a form of abstraction), while the detailed state information is kept on the stack. These techniques allow the detection of subtle bugs which can not be discovered by classical model checking (using e.g. breadth first search) or by state-less model checking [8]. While these techniques use abstractions in an ad-hoc manner, our work contributes the automated generation and refinement of abstractions.

Dataflow and type-based analyzes have been used to check safety properties of software (e.g. [20]). Unlike our work, these techniques analyze over-approximations of system behavior and may generate false positive results due to infeasible paths.

Layout The rest of the paper is organized as follows. Section 2 shows an example illustrating our approach. Section 3 gives background information. Section 4 describes the main algorithm for performing concrete model checking with abstract matching and refinement. Section 5 discusses correctness and termination for the algorithm. Section 6 proposes extensions to the main algorithm. Section 7 illustrates applications of the approach and Section 8 concludes the paper.

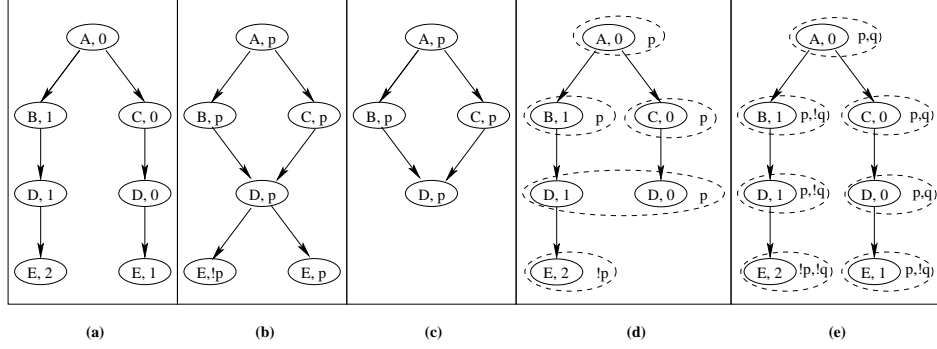


Fig. 1. (a) Concrete system (b) *May* abstraction using predicate $p = x < 2$ (c) *Must* abstraction using p (d) Concrete search with abstract matching using p (e) Concrete search with abstract matching using predicates p and $q = x < 1$.

2 Example

The example in Fig. 1 illustrates some of the main characteristics of our approach. Fig. 1 (a) shows the state space of a concrete system that has only one variable x ; states are labelled with the program counter (e.g. A, B, C ...) and the concrete value of x . Fig. 1 (b) shows the abstract system induced by the *may* transitions for predicate $p = x < 2$. Fig. 1 (c) shows the abstract system induced by the *must* transitions for predicate p .

Fig. 1 (d) shows the state space that is explored using our proposed approach, for an abstraction specified by predicate p . Dotted circles denote the abstract states which are stored, and used for matching, during the concrete execution of the system. The approach explores only the *feasible* behavior of the concrete system, following transitions that correspond to both *may* and *must* transitions, but it might miss behavior due to abstract matching. For example, the state labelled by $(E, 1)$ is not explored (assuming a breadth-first search), since $(D, 0)$ was matched with $(D, 1)$ - both have the same program counter and both satisfy p . Notice that, with respect to reachable states, the produced structure is a better under-approximation than the *must* abstraction. Fig. 1 (e) illustrates concrete execution with abstract matching, after a refinement step, which introduced a new predicate $q = x < 1$. The resulting structure is an exact abstraction of the concrete system.

3 Background

Program Model To make the presentation simple, we use as a specification language a guarded commands language over integer variables. Most of the results extend directly to more sophisticated programming languages. Let V be a finite set of integer variables. Expressions over V are defined using standard boolean ($=, <, >$) and binary ($+, -, \cdot, \dots$) operations. A *model* is a tuple $M = (V, T)$.

$T = \{t_1, \dots, t_n\}$ is a finite set of transitions, where $t_i = (g_i(\mathbf{x}) \mapsto \mathbf{x} := e_i(\mathbf{x}))$. $g_i(\mathbf{x})$ is a guard and $e_i(\mathbf{x})$ are assignments to the variables represented by tuple \mathbf{x} ; throughout the paper, we write this as a sequence of assignments.

Semantics As a semantics of a model we use transition systems. A *transition system* over a finite set of atomic propositions AP is a tuple (S, R, s_0, L) where S is a (possibly infinite) set of states, $R = \{\xrightarrow{i}\}$ is a finite set of deterministic transition relations: $\xrightarrow{i} \subseteq S \times S$, s_0 is an initial state, and $L : S \rightarrow 2^{AP}$ is a labelling function. State s is *reachable* if it is reachable from the initial state via zero or more transitions, i.e. $s_0 \rightarrow^* s$. The set of *reachable labellings* RL is $\{L(s) \mid \exists s \in S : s_0 \rightarrow^* s\}$. The *concrete semantics* of model M is the transition system $C(M) = (S, \{\xrightarrow{i}\}, s_0, L)$ over AP , where:

- $S = 2^{V \rightarrow \mathbb{Z}}$, i.e., states are valuations of variables
- $s \xrightarrow{i} s' \Leftrightarrow s \models g_i \wedge s' = u_i(s)$, where the semantics of an expression and update is as usual, expressions are functions $\models_e : (V \rightarrow \mathbb{Z}) \rightarrow \{\text{true}, \text{false}\}$, written as $e \models s$; updates are functions $u_i : (V \rightarrow \mathbb{Z}) \rightarrow (V \rightarrow \mathbb{Z})$.
- s_0 is the zero valuation ($\forall v \in V : s_0(v) = 0$)
- $L(s) = \{p \in AP \mid s \models p\}$

Weakest precondition The *weakest precondition* of a set of states X with respect to transition i is $wp(X, i) = \{s \mid s \xrightarrow{i} s' \Rightarrow s' \in X\}$. If the set of states X is characterized by a predicate ϕ , then the weakest precondition with respect to transition i can be expressed as $wp(\phi, i) = (g_i \Rightarrow \phi[e_i(\mathbf{x})/\mathbf{x}])$.

Predicate abstraction Predicate abstraction is a special instance of the framework of abstract interpretation [4] that maps a (potentially infinite state) transition system into a finite state transition system via a set of predicates $\Phi = \{\phi_1, \dots, \phi_n\}$ over the program variables. Let \mathbb{B}_n be a set of bitvectors of length n . We define abstraction function $\alpha_\Phi : S \rightarrow \mathbb{B}_n$, such that $\alpha_\Phi(s)$ is a bitvector $b_1 b_2 \dots b_n$ such that $b_i = 1 \Leftrightarrow s \models \phi_i$. Let Φ_s be the set of all abstraction predicates that evaluate to *true* for a given state s , i.e., $\Phi_s = \{\phi \in \Phi \mid s \models \phi\}$. For succinctness we sometimes write $\alpha_\Phi(s)$ (or just $\alpha(s)$) to denote $\bigwedge_{\phi \in \Phi_s} \phi$.

We also give here the definitions of *may* and *must* abstract transitions. Although not necessary for formalizing our algorithm, these definitions clarify the comparison with related work. For two abstract states (bit vectors) a_1 and a_2 :

- $\xrightarrow{must} : a_1 \xrightarrow{must} a_2$ iff for all concrete states s_1 such that $\alpha(s_1) = a_1$, there exists concrete state s_2 such that $\alpha(s_2) = a_2$ and $s_1 \xrightarrow{i} s_2$
- $\xrightarrow{may} : a_1 \xrightarrow{may} a_2$ iff there exists concrete state s_1 such that $\alpha(s_1) = a_1$ and there exists concrete state s_2 such that $\alpha(s_2) = a_2$, such that $s_1 \xrightarrow{i} s_2$

Algorithms for computing abstractions using over-approximation based predicate abstraction are given in e.g. [1, 9] (they compute *may* abstract transitions automatically, with the help of a theorem prover). In the worst case, these algorithms make $2^n \times n \times 2$ calls to the theorem prover for each program transition. Note that our approach does not require the computation of abstract transitions, since it executes directly the concrete transitions.

Bisimulation A symmetric relation $R \subseteq S \times S$ is a bisimulation relation iff for all $(s, s') \in R$:

- $L(s) = L(s')$
- For every $s' \xrightarrow{i} s'_1$ there exists $s \xrightarrow{i} s_1$ such that $R(s_1, s'_1)$.

The bisimulation is the largest bisimulation relation, denoted \sim . Two transition systems are bisimilar if their initial states are bisimilar.

As \sim is an equivalence relation, it induces a *quotient* transition system whose states are equivalence classes with respect to \sim and there is a transition between two equivalence classes A and B if $\exists s_1 \in A$ and $\exists s_2 \in B$ such that $s_1 \xrightarrow{i} s_2$.

4 Concrete Model Checking with Abstract Matching

Algorithm Fig. 2 shows the reachability procedure that performs model checking with abstract matching (α SEARCH). It is basically concrete state space exploration with matching on abstract states; the main modification with respect to classical state space search is that we store $\alpha(s)$ instead of s . The procedure uses the following data structures:

- *States* is a set of abstract states visited so far
- *Transitions* is a set of abstract transitions visited so far
- *Wait* is a set of concrete states to be explored

The procedure performs validity checking, using a theorem prover, to determine whether the abstraction is *exact* with respect to each explored transition – see discussion below. Φ_{new} maintains the list of abstraction predicates. The procedure returns the computed structure and a set of new predicates, that are used for refinement.

Fig. 3 gives the iterative refinement algorithm for checking whether M can reach an error state described by φ . At each iteration of the loop, the algorithm invokes procedure α SEARCH to analyze an under-approximation of the system, which either violates the property, it is proved to be correct (if the abstraction is found to be exact with respect to all transitions), or it needs to be refined.

Checking for Exact Abstraction and Refinement We say that an abstraction function α is *exact* with respect to transition $s \xrightarrow{i} s'$ iff $\forall s_1$ such that $\alpha(s) = \alpha(s_1)$, $\exists s'_1$, such that $\alpha(s'_1) = \alpha(s')$ and $s_1 \xrightarrow{i} s'_1$. In other words, α is *exact* with respect to $s \xrightarrow{i} s'$ iff $\alpha(s) \xrightarrow{i}_{must} \alpha(s')$. This definition is also related to the notion of *completeness* in abstract interpretation (see e.g. [7]), which states that no loss of precision is introduced by the abstraction.

Checking that the abstraction is *exact* with respect to concrete transition $s \xrightarrow{i} s'$ is equivalent to checking that $\alpha_\Phi(s) \Rightarrow wp(\alpha_\Phi(s'), i)$ is valid. This formula is equivalent to $\alpha_\Phi(s) \Rightarrow \alpha_\Phi(s')[e_i(\mathbf{x})/\mathbf{x}]$ when $s \models g_i$. Checking the validity for these formulas is in general undecidable. As is customary, if the

```

proc  $\alpha\text{SEARCH}(M, \Phi)$ 
   $\Phi_{new} = \Phi$ ; add  $s_0$  to Wait; add  $\alpha_\Phi(s_0)$  to States
  while Wait  $\neq \emptyset$  do
    get  $s$  from Wait
     $L(\alpha_\Phi(s)) = \{a \in AP \mid s \models a\}$ 
    foreach  $i$  from 1 to  $n$  do
      if  $s \models g_i$  then
        if  $\alpha_\Phi(s) \Rightarrow g_i$  is not valid
          then add  $g_i$  to  $\Phi_{new}$  fi
         $s' = u_i(s)$ 
        if  $\alpha_\Phi(s) \Rightarrow \alpha_\Phi(s')[e_i(\mathbf{x})/\mathbf{x}]$  is not valid
          then add predicates in  $\alpha_\Phi(s')[e_i(\mathbf{x})/\mathbf{x}]$  to  $\Phi_{new}$  fi
        if  $\alpha_\Phi(s') \notin \text{States}$  then
          add  $s'$  to Wait
          add  $\alpha_\Phi(s')$  to States
        fi
        add  $(\alpha_\Phi(s), i, \alpha_\Phi(s'))$  to Transitions
      else
        if  $\alpha_\Phi(s) \Rightarrow \neg g_i$  is not valid
          then add  $g_i$  to  $\Phi_{new}$  fi
      fi
    od
   $A = (\text{States}, \text{Transitions}, \alpha_\Phi(s_0), L)$ 
  return  $(A, \Phi_{new})$ 
end

```

Fig. 2. Search procedure with checking for exact abstraction

theorem prover can not decide the validity of a formula, we assume that it is not valid. This may cause some unnecessary refinement, but it keeps the correctness of the approach. If the abstraction can not be proved to be exact with respect to some transition, then the new predicates from the failed formula are added to the set of abstraction predicates. Intuitively, these predicates will be useful for proving exactness in the next iteration.

When updating Φ_{new} for refinement, it is sufficient to add only those conjuncts of $\alpha_\Phi(s')[e_i(\mathbf{x})/\mathbf{x}]$ for which we cannot prove validity. Note that it is not necessary to perform the validity check during model checking. In fact, our implementation collects all the theorem prover queries and it runs the theorem prover off-line. In this way, we also make sure that the theorem prover is not called twice for the same query.

5 Correctness and Termination

In this section we discuss the properties of the refinement algorithm. We state only the main theorems, technical lemmas and proofs are in the appendix (due to space limitations). First, we show that the set $RL(\alpha\text{SEARCH}(M, \Phi))$ of reachable

```

proc REFINEMENTSEARCH( $M, \varphi$ )
   $i = 1; \Phi_i = \emptyset$ 
  while true do
     $(A_i, \Phi_{i+1}) = \alpha\text{SEARCH}(M, \Phi_i)$ 
    if  $\varphi$  is reachable in  $A_i$  then return counter-example fi
    if  $\Phi_{i+1} = \Phi_i$  then return unreachable fi
     $i = i + 1$ 
  od
end

```

Fig. 3. Iterative refinement algorithm

labellings computed by the algorithm REFINEMENTSEARCH is a subset of the reachable labelings of the system under analysis.

Theorem 1. *Let $AP \subseteq \Phi$. Then $RL(\alpha\text{SEARCH}(M, \Phi)) \subseteq RL(C(M))$.*

Moreover, it holds that $RL(\alpha\text{SEARCH}(M, \Phi))$ is a superset of the reachable labellings in the *must* abstraction (see Lemma 2 in Appendix), hence it is a better approximation.

We now show that, if the iterative algorithm terminates then the result is correct and moreover, if the error state is unreachable, the output structure is bisimilar to the system under analysis:

Theorem 2. *If REFINEMENTSEARCH(M, φ) terminates then:*

- *If it returns a counter-example, then it is a real error.*
- *If it returns 'unreachable', then the error state is indeed unreachable in M and moreover the computed structure is bisimilar to $C(M)$.*

In general, the proposed algorithm might not terminate (because of the halting problem). However, the algorithm is guaranteed to eventually find all the reachable labellings of the concrete program, although it might not be able to detect that (to decide termination). Moreover, if the (reachable part of the) system under analysis has a finite bisimulation quotient, then the algorithm will eventually produce a finite bisimilar structure.

Theorem 3. *Let the αSEARCH use breadth-first search order and let $A_1, A_2 \dots$ be a sequence of transition systems generated during iterative refinement performed by REFINEMENTSEARCH. Then*

- *There exists i such that $RL(A_i) = RL(C(M))$.*
- *If the reachable part of the bisimulation quotient is finite, then there exists i such that $A_i \sim C(M)$.*

Discussion The search order used in αSEARCH (depth-first or breadth-first) influences the size of the generated structure, the newly computed predicates, and even the number of iterations of the main algorithm. If there are two states

s_1 and s_2 such that $\alpha_\Phi(s_1) = \alpha_\Phi(s_2)$ but $s_1 \not\sim s_2$ then, depending on whether s_1 or s_2 is visited first, different parts of the transition system will be explored.

Also note that the refinement algorithm is non-monotone, i.e. a labelling which is reachable in one iteration may not be reachable in the next iteration. This is in conformance to the well known result that the set of *must* transitions is not generally monotonically non-decreasing when predicates are added to refine an abstract system. However, we should note that the algorithm is guaranteed to converge to the correct answer.

We should also note that the proposed iterative algorithm is not guaranteed to terminate even for a finite state program. This situation is illustrated by the following example (the property we are checking is that $pc = 2$ is unreachable).

$$\begin{aligned} pc = 0 &\mapsto x := 0, y := 0, pc := 1 \\ pc = 1 \wedge y \geq 0 &\mapsto y := y + x \\ pc = 1 \wedge y < 0 &\mapsto pc := 2 \end{aligned}$$

Although the program is finite state (and therefore the problem can be easily solved with classical explicit model checking), it is quite difficult to solve using abstraction refinement techniques. The iterative algorithm will not terminate on this example: it will keep adding predicates $y \geq 0, y + x \geq 0, y + 2x \geq 0, \dots$. Note that, in accordance with Theorem 3, it will eventually produce a bisimilar structure. However, the algorithm will not be able to detect termination, and it will keep refining indefinitely. The reason is that the algorithm keeps adding predicates that refine the unreachable part of the system under analysis. Also note that the same problem will occur with over-approximation based abstraction techniques that use refinement based on weakest precondition calculations [2, 16]. Those techniques will introduce the same predicates.

To solve this problem, we propose to use the following heuristic. If there is a transition for which we cannot prove that the abstraction is exact in several subsequent iterations of the algorithm, then we add predicates describing the concrete state; i.e., in our example we would add predicates $x = 0; y = 0$. The abstraction will eventually become exact with respect to each transition. And since the number of reachable transitions is finite, the algorithm must terminate.

Lemma 1. *If $C(M)$ is finite state then the modified algorithm terminates.*

6 Extensions

Lightweight Approach As mentioned, the under-approximation and refinement approach can be used in a lightweight but systematic manner, without using a theorem prover for validity checking. Specifically, for each explored transition t_i refinement adds the new predicates from $\alpha_\Phi(s')[e_i(\mathbf{x})/\mathbf{x}]$, regardless of the fact that the abstraction is exact with respect to transition t_i . This approach may result in unnecessary refinement. A similar refinement procedure was used in [16] for automated over-approximation predicate abstraction.

We are also considering several heuristics for generating new abstraction predicates. For example, it is customary to add the predicates that appear in

the guards and in the property to be checked. One could also add predicates generated dynamically, using tools like Daikon [6], or predicates from known invariants of the system (generated using static analysis techniques).

In order to extend the applicability of the proposed technique to the analysis of full-fledged programming languages, we are investigating shape abstractions [18], that record information about the shape of the program heap, to be used in conjunction with the abstraction predicates. Section 7 shows an example use of shape abstraction for the analysis of Java programs.

Transition Dependent Predicates The predicates that are generated after the validity check for one transition are used ‘globally’ at the next iteration. This may cause unnecessary refinement — the new predicates may distinguish states which do not need to be distinguished. To avoid this, we could use ‘transition dependent’ predicates. The idea is to associate the abstraction predicates with the program counter corresponding to the transition that generated them. New predicates are then added only to the set of the respective program counter. However, with this approach, it may take longer before predicates are ‘propagated’ to all the locations where they are needed, i.e. more iterations are needed before an error is detected or an exact abstraction is found. We need to further investigate these issues. Similar ideas are presented in [3], with the purpose of minimizing the number of predicates for over-approximation based abstraction.

7 Applications

We have implemented our approach for the guarded command language. Our implementation is done in the language Perl and it uses the Simplify theorem prover [5]. We discuss the application of our implementation for two concurrent programs: property verification for the Bakery mutual exclusion protocol and error detection in RAX (Remote Agent Experiment), a component extracted from an embedded spacecraft-control application.

These preliminary experiments show the merits of our approach. Of course, much more experimentation is necessary to really assess the practical benefits of the proposed technique and a lot more engineering is required to apply it to real programming languages. We are currently doing an implementation in the Java PathFinder (JPF) model checking framework [21] for the analysis of Java programs (we use the Omega library as a decision procedure). We briefly discuss at the end of this section the use of our approach for test-case generation for Java container implementations.

The Bakery Mutual Exclusion Protocol We have analyzed several versions of the Bakery mutual exclusion protocol (for two and more processes). These versions are infinite state but they have a finite bisimulation quotient. The guarded command representation for a simplified version of the protocol is given in Fig. 4.

The mutual exclusion property is encoded as “ $pc_1 = 3 \wedge pc_2 = 3$ is unreachable”. We used our tool to successfully prove that the property holds. Fig. 5 gives the intermediate results of the analysis. For each iteration, we report the number

(Process 1)	(Process 2)
$pc_1 = 0 \mapsto x := y, pc_1 := 1$	$pc_2 = 0 \mapsto y := x, pc_2 := 1$
$pc_1 = 1 \mapsto x := x + 1, pc_1 := 2$	$pc_2 = 1 \mapsto y := y + 1, pc_2 := 2$
$pc_1 = 2 \wedge x \leq y \mapsto pc_1 := 3$	$pc_2 = 2 \wedge y < x \mapsto pc_2 := 3$
$pc_1 = 3 \mapsto pc_1 := 0$	$pc_2 = 3 \mapsto pc_2 := 0$

Fig. 4. Bakery example

Iteration	Concrete states	Abstract states	New predicates
1	17	11	$x \leq y$
2	18	12	$x + 1 \leq y, x \leq y + 1, y \geq 0$
3	26	19	$x + 2 \leq y, y \geq 1, x \leq 1$
4	44	32	$y \leq 1, x \leq 0, y \geq 2$
5	48	36	-

Fig. 5. Bakery example: intermediate results of the refinement algorithm

of generated concrete states, the number of stored abstract states and the newly generated predicates. Note that we never abstract the program counter. The reported results are for the breadth-first search order. For the depth-first search order the algorithm requires only 4 iterations (see the discussion in Section 5). The algorithm proceeds in similar way for the full version of the protocol.

RAX The RAX example (illustrated in Fig.6) is derived from the software used within the NASA Deep Space 1 Remote Agent experiment, which deadlocked during flight [22]. We encoded the deadlock check as “ $pc_1 = 4 \wedge pc_2 = 5 \wedge w_1 = 1 \wedge w_2 = 1$ is unreachable”. The error is found after one iteration, for breadth-first search order; the reported counter-example has 8 steps. For depth-first search order, the algorithm needs one more iteration to find the error, using the predicates that appear in the guards $c_1 = e_1$ and $c_2 = e_2$. Note that the state space of the program is unbounded, as the program keeps incrementing the counters e_1 and e_2 , when $pc_2 = 2$ and $pc_1 = 6$, respectively. We also ran our algorithm to see if it converges to a finite bisimulation quotient. Interestingly, the algorithm does not terminate as it keeps adding predicates in a way similar to the example presented in Section 5 (we report the results in the appendix).

Note that the application of over-approximation based predicate abstraction to a Java version of RAX is described in detail in [22]. In that work, four different predicates were used to produce an abstract model that is bisimilar to the original program. In contrast, the work presented here allowed more aggressive abstraction to recover feasible counter-examples.

In general, we believe that the technique presented here is complementary to over-approximation abstraction methods and it can be used in conjunction with such methods, as an efficient way of discovering feasible counter-examples. We view the integration of the two approaches as an interesting research topic for the future. Our technique explores transitions that are guaranteed to be feasible in the state space bounded by the abstraction predicates. In contrast, the over-approximation based methods may also explore transitions that are

(Process 1)	(Process 2)
$pc_1 = 1 \mapsto c_1 := 0, pc_1 := 2$	$pc_2 = 1 \mapsto c_2 := 0, pc_2 := 2$
$pc_1 = 2 \wedge c_1 = e_1 \mapsto pc_1 := 3$	$pc_2 = 2 \mapsto e_1 := e_1 + 1, w_1 := 0, pc_2 := 3$
$pc_1 = 3 \mapsto w_1 := 1, pc_1 := 4$	$pc_2 = 3 \wedge c_2 = e_2 \mapsto pc_2 := 4$
$pc_1 = 4 \wedge w_1 = 0 \mapsto pc_1 := 5$	$pc_2 = 4 \mapsto w_2 := 1, pc_2 := 5$
$pc_1 = 2 \wedge c_1 \neq e_1 \mapsto pc_1 := 5$	$pc_2 = 5 \wedge w_2 = 0 \mapsto pc_2 := 6$
$pc_1 = 5 \mapsto c_1 := e_1, pc_1 := 6$	$pc_2 = 3 \wedge c_2 \neq e_2 \mapsto pc_2 := 6$
$pc_1 = 6 \mapsto e_2 := e_2 + 1, w_2 := 0, pc_1 := 2$	$pc_2 = 6 \mapsto c_2 := e_2, pc_2 := 2$

Fig. 6. RAX example

spurious and therefore could require additional refinement before reporting a real counter-example. Hence, our technique can potentially finish in fewer iterations and it can use fewer predicates (which enable more state space reduction), while retaining the model checker's capability of finding real bugs.

Testing We have used our preliminary implementation in the JPF model checker to perform test case generation to achieve code coverage for Java container classes (tree-map, linked-list, fibonacci-heap). Test cases are sequences of API calls, i.e. method calls that *add* and *remove* elements in a container, to obtain for example, branch coverage. The model checker analyzes all sequences of API calls up to a predefined sequence size and uses abstraction mappings to match states between API calls to avoid the generation of redundant tests. We used an abstraction recording the shape of the containers augmented with different predicate abstractions on the data fields from each container element - two states are matched if they represent containers that have the same shape and valuation for the abstraction predicates. The behavioral coverage obtained in this fashion is highly dependent on the different abstractions that are used. Therefore we believe that the capability of generating and refining the abstractions automatically is crucial for achieving good coverage. Although the work presented here is only a first step towards this goal (the JPF implementation does not yet allow automated refinement), we obtained better behavioral coverage than with exhaustive model checking. In fact, for some of the examples, exhaustive analysis runs out of memory even before generating tests that cover all the reachable branches in the code.

8 Conclusions and Future Work

We presented a novel model checking algorithm based on refinement of under-approximations, which effectively preserves the defect detection ability of model checking in the presence of powerful abstractions. The under-approximation is obtained by traversing the concrete transition system and performing the state matching on abstract states computed by predicate abstraction. The refinement is done by checking exactness of abstractions with the use of a theorem prover. We illustrated the application of the algorithm for checking safety properties of concurrent programs and for testing container implementations. In the future, we plan to extend the algorithm to checking liveness properties. We also plan to do an extensive evaluation of our approach on real systems.

References

1. T. Ball, A. Podelski, and S. Rajamani. Boolean and Cartesian Abstractions for Model Checking C Programs. In *Proc. of TACAS'01*, April 2001.
2. S. Chaki, E. Clarke, A. Groce, S. Jha, and H. Veith. Modular Verification of Software Components in C. *ACM Trans. Computer Systems*, 30(6):388–402, 2004.
3. S. Chaki, E. Clarke, A. Groce, and O. Strichman. Predicate abstraction with minimum predicates. In *12th CHARME*, 2003.
4. P. Cousot and R. Cousot. Abstract Interpretation Frameworks. *Journal of Logic and Computation*, 4(2):511–547, August 1992.
5. D. L. Detlefs, K. R. M. Leino, and J. B. Saxe G. Nelson. Extended Static Checking. *Research Report 159, Compaq Systems Research Center*, 1998.
6. M. D. Ernst, A. Czeisler, W. G. Griswold, and D. Notkin. Quickly detecting relevant program invariants. In *Proc. of 22nd ICSE*, 2000.
7. R. Giacobazzi and E. Quintarelli. Incompleteness, counterexamples and refinements in abstract model checking. In *Proc. 8th SAS*, 2001.
8. P. Godefroid. Software Model Checking: the Verisoft Approach. *Formal Methods in Systems Design (to appear)*.
9. S. Graf and H. Saidi. Construction of abstract state graphs with PVS. In *Proc. CAV*, 1997.
10. W. Grieskamp, Y. Gurevich, W. Schulte, and M. Veanes. Generating finite state machines from abstract state machines. In *Proc. ISSTA*, July 2002.
11. O. Grumberg, F. Lerda, O. Strichman, and M. Theobald. Proof-guided underapproximation-widening for multi-process systems. In *Proc. 32nd POPL*, 2005.
12. T. A. Henzinger, R. Jhala, R. Majumdar, and Gregoire Sutre. Lazy abstraction. In *Proc. POPL*, pages pp. 58–70. ACM Press, 2002.
13. G. J. Holzmann and R. Joshi. Model-driven software verification. In *Proc. 11th SPIN Workshop*, Barcelona, Spain, 2004.
14. D. Lee and M. Yannakakis. Online minimization of transition systems. In *Proc. 24th ACM Symposium on Theory of Computing*, 1992.
15. M. Musuvathi, D. Y. W. Park, A. Chou, D. R. Engler, and D. L. Dill. CMC: A pragmatic approach to model checking real code. In *Proc. 5th Symposium on Operating Systems Design and Implementation*, 2002.
16. K. S. Namjoshi and R. P. Kurshan. Syntactic program transformations for automatic abstraction. In *Proc. 12th CAV*, 2000.
17. C. S. Pasareanu, M. B. Dwyer, and W. Visser. Finding Feasible Abstract Counterexamples. *STTT*, 5(1):34–48, November 2003.
18. S. Sagiv, T. W. Reps, and R. Wilhelm. Parametric shape analysis via 3-valued logic. *TOPLAS*, 2002.
19. S. Shoham and O. Grumberg. Monotonic abstraction-refinement for CTL. In *Proc. TACAS*, Barcelona, Spain, 2004.
20. A. Venet and G. Brat. Precise and efficient static array bound checking for large embedded c programs. In *Proc. PLDI*, 2004.
21. W. Visser, K. Havelund, G. Brat, S. J. Park, and F. Lerda. Model Checking Programs. *ASE Journal*, 10(2), April 2003.
22. W. Visser, S. Park, and J. Penix. Applying predicate abstraction to model check object-oriented programs. In *3rd ACM SIGSOFT Workshop on Formal Methods in Software Practice*, 2000.
23. T. Xie, D. Marinov, and D. Notkin. Rostra: A framework for detecting redundant object-oriented unit tests. In *Proc. 19th ASE*, Linz, 2004.

A Appendix

We provide here the proofs of several technical lemmas. The theorems in Section 5 are direct consequences of these lemmas. We also give the results of running the algorithm on RAX with the goal of computing a finite bisimilar system.

A.1 Technical Lemmas and Proofs

We use the following notation: a state s was *visited* during the search if it was inserted into *Wait*, that it is *considered* if it is generated as a successor of some state in the **foreach** loop. State s_1 is *matched* to state s_2 if the check $\alpha_\Phi(s_1) \notin \text{States}$ fails because $\alpha_\Phi(s_1) = \alpha_\Phi(s_2)$ and s_2 was inserted into *States* before. We abuse the notation and sometimes we let $\alpha\text{SEARCH}(M, \Phi)$ denote just the structure A computed by the algorithm and not the tuple (A, Φ_{new}) . We say that transition t_i is exact if α_Φ is exact with respect to it.

For the proofs, we need to refine the definition of bisimulation. A symmetric relation $R \subset S \times S$ is a k -bisimulation relation iff:

- for all $(s, s') \in R : L(s) = L(s')$
- if $k > 0$ then there exists $(k-1)$ -bisimulation relation R' such that for all $(s, s') \in R : (\forall s \xrightarrow{i} s_1 \Rightarrow \exists s' \xrightarrow{i} s'_1 \wedge (s, s') \in R')$

The k -bisimulation is the largest k -bisimulation relation, denoted \sim_k .

Lemma 2. *If a state s is reachable in $C(M)$ via exact transitions with respect to Φ , then there exists s' such that s' is visited during the $\alpha\text{SEARCH}(M, \Phi)$ and $\alpha_\Phi(s) = \alpha_\Phi(s')$.*

Proof: By induction with respect to the number of exact transitions from the initial state. Basic step ($k = 0$) is trivial. For the induction step, suppose that state s is reachable via sequence of exact transitions: $s_0 \xrightarrow{i_0} \dots \xrightarrow{i_{k-1}} s_k \xrightarrow{i_k} s_{k+1} = s$. By induction hypothesis there exists s'_k such that s'_k is visited and $\alpha_\Phi(s'_k) = \alpha_\Phi(s_k)$. Because the abstraction is exact with respect to $s_k \xrightarrow{i_k} s$, there must be s' such that $s'_k \xrightarrow{i_k} s'$ and $\alpha_\Phi(s') = \alpha_\Phi(s)$. This successor s' was considered during the visit of s'_k . There are two cases to be considered.

- s' was added to *Wait* and later visited
- s' was matched to a previously visited s'' such that $\alpha_\Phi(s') = \alpha_\Phi(s'')$

In any case we get that some state with the same $\alpha_\Phi(s)$ was previously visited.

Lemma 3. *Let $AP \subseteq \Phi$. Then $RL(\alpha\text{SEARCH}(M, \Phi)) \subseteq RL(C(M))$*

Proof: It is easy to verify that the following is an invariant of the search: '*States*' is a subset of reachable states in $C(M)$.

Lemma 4. *Let $AP \subseteq \Phi$. If for all reachable states $s_1, s_2 : \alpha_\Phi(s_1) = \alpha_\Phi(s_2) \Rightarrow s_1 \sim s_2$, then $\alpha\text{SEARCH}(M, \Phi) \sim C(M)$.*

Proof: Consider relation R defined as: $s_1 R s_2$ iff $s_1 = s_2$ or s_1 was matched to s_2 . It is easy to verify that R is a bisimulation relation between $\alpha\text{SEARCH}(M, \Phi)$ and $C(M)$.

Lemma 5. *Let $(A, \Phi_{new}) = \alpha\text{SEARCH}(M, \Phi)$. If $\Phi_{new} = \Phi$, then $A \sim C(M)$.*

Proof: Due to Lemma 4 it is sufficient to show that α_Φ induces a bisimulation relation on the reachable part of the transition system. We proceed by contradiction. Suppose that there exists reachable states s_1, s_2 such that $\alpha_\Phi(s_1) = \alpha_\Phi(s_2)$ and there exists $s_1 \xrightarrow{i} s'_1$ such that $s_2 \xrightarrow{i} s'_2$ and $\alpha(s'_1) \neq \alpha(s'_2)$. This means that α_Φ is not exact with respect to $s_1 \xrightarrow{i} s'_1$.

Let us consider such an s_1 which is closest to initial state. This state must be reachable only via exact transitions. According to Lemma 2 some state s such that $\alpha(s) = \alpha(s_1)$ was visited during the search. During the visit of the state s we check whether the abstraction is exact for transition $s \xrightarrow{i} s'$. Since it is not exact, some of the implications is not valid and therefore Φ_{new} is updated and thus $\Phi \neq \Phi_{new}$ and that is the contradiction.

Lemma 6. *Let $\{A_i\}_{i=0}^\infty$ be a sequence of transition systems generated during an infinite run of REFINEMENTSEARCH and $\text{Inf}_M = \{s \mid \text{there exists infinitely many } i \text{ such that } s \in A_i\}$. If $s \not\sim_k s'$ and $s \in \text{Inf}_M$ then there exists i such that $\alpha_{\Phi_i}(s) \neq \alpha_{\Phi_i}(s')$.*

Proof: We prove by induction with respect to k where k is the smallest number such that $s \not\sim_k s'$. Basic step: for $k = 0$ it means that $L(s) \neq L(s')$ and therefore $\alpha_{\Phi_1}(s) \neq \alpha_{\Phi_1}(s')$. Induction step ($k + 1$): Let s_1, s'_1 be such that $s \xrightarrow{i} s_1, s' \xrightarrow{i} s'_1$ and $s_1 \not\sim_k s'_1$. Then one of the following must hold:

1. State $s_1 \in \text{Inf}_M$. Then we can apply induction hypothesis, i.e., eventually $\alpha_{\Phi_i}(s_1) \neq \alpha_{\Phi_i}(s'_1)$.
2. State s_1 is matched to some state in infinitely many runs of αSEARCH . Since we use breadth-first order, there are only finitely many states to which it can be matched. Therefore, there exists a state s_2 such that s_1 is matched to s_2 in infinitely many runs of αSEARCH . From induction hypothesis we get that $s_2 \not\sim_k s'_1$ and that eventually $\alpha_{\Phi_i}(s_2) \neq \alpha_{\Phi_i}(s'_1)$. Because s_1 is matched to s_2 infinitely often we eventually get also $\alpha_{\Phi_j}(s_1) \neq \alpha_{\Phi_j}(s'_1)$ for some $j \geq i$.

This means that α_{Φ_i} is not exact with respect to $s \xrightarrow{i} s_1$, therefore $wp(\alpha_{\Phi_i}(s_1), t_i)$ will be included in Φ_{i+1} and therefore $\alpha_{\Phi_{i+1}}(s_1) \neq \alpha_{\Phi_{i+1}}(s'_1)$.

Lemma 7. *For each reachable bisimulation class B there exists a state $s \in B$ such that s is visited by REFINEMENTSEARCH .*

Proof: By induction with respect to the length of the shortest path in which state from B is reachable. Base step is obvious. Induction step: let state from B be reachable via path s_0, \dots, s_k, s_{k+1} . By induction hypothesis some state $s' \sim s_k$ is reached during the search. And with the use of Lemma 6 we get that some $s'' \sim s_{k+1}$ is reached.

Iteration	Concrete states	Abstract states	New predicates
1	56	35	$c_1 = e_1, c_2 = e_2$
2	68	44	$e_1 = 0, e_2 = 0$
3	100	65	$e_1 = -1, e_2 = -1$
4	100	65	$e_1 = -2, e_2 = -2$
5	100	65	...

Fig. 7. RAX example: intermediate results of the refinement algorithm

Lemma 8. *Let $\{A_i\}_{i=0}^{\infty}$ be a sequence of transition systems generated during an infinite run of $\text{REFINEMENTSEARCH}(M, \varphi)$. There exists i such that $RL(A_i) = RL(C(M))$.*

Proof: For each $l \in RL(C(M))$ we choose some bisimulation class B such that $s \in B \Rightarrow L(s) = l$. In this way we obtain a finite set of bisimulation classes $\{B_1, \dots, B_k\}$ which “covers” $RL(C(M))$ (note that $RL(C(M))$ is finite because AP is finite). Now we show that there exists an iteration in which at least one state from each of these classes is visited. This is done similarly to the proof of Lemma 7.

Lemma 9. *Let $\{A_i\}_{i=0}^{\infty}$ be a sequence of transition systems generated during an infinite run of REFINEMENTSEARCH . If the reachable part of bisimulation quotient is finite, then there exists i such that $A_i \sim C(M)$.*

Proof: By contradiction. Suppose that $\forall i : A_i \not\sim C(M)$. From Lemma 4 we get that there exists reachable s, s' such that $\forall i : \alpha_{\phi_i}(s) = \alpha_{\phi_i}(s')$ and $s \not\sim s'$. We show (similarly to the proof of Lemma 2) that there exists such s which is visited infinitely often. From Lemma 6 we get that eventually $\alpha_{\phi_i}(s) \neq \alpha_{\phi_i}(s')$ which is the contradiction.

A.2 RAX example: intermediate results of the refinement algorithm

As mentioned, the RAX example has a deadlock, which is detected by our algorithm after the first iteration. As the program is infinite state, we also tried to use the algorithm to compute a finite bisimilar structure. Interestingly, the algorithm does not terminate for the RAX example (although it has a finite bisimulation quotient). The results are shown in Figure 7 (breadth-first search order). However, if we assume that the counters in the program are non-negative, i.e if we introduce two new predicates, $e_1 \geq 0, e_2 \geq 0$, then the algorithm terminates after three iterations.